

TDMA BACKBONE FOR A MOBILE NETWORK

FIELD OF THE INVENTION

[0001] The present invention relates to communication networks and, more particularly, to mobile communication networks.

BACKGROUND OF THE INVENTION

[0002] The effectiveness of a communication network depends essentially on its network architecture, components and operating conditions. Antennas are critical components in a mobile communication network. A properly chosen antenna can significantly enhance the performance of a mobile communication network.

[0003] A phased array antenna (PAA) employs spatially narrow radio frequency (RF) beams to communicate between the network nodes. Electronic beam steering allows the PAA beams to be retargeted instantly after a gap of minimum set-up time for loading the PAA shift registers. In a synchronous link approach the transmit beam of a first PAA is pointed towards the receive beam of a second PAA. Similarly, the transmit beam of the second PAA is pointed towards the receive beam of the first PAA. In the synchronous approach, each node having PAAs can form no more links than the number of beams on its PAAs. Consequently, the network's performance is limited by the number of beams on the PAAs.

[0004] Simulations of typical PAA based networks show that for a given spatial distribution of the nodes, the number of links actually available is much smaller than the maximum number of available links. This results in the average packet route going through several node hops, which is very inefficient. Adding more PAAs imposes increased demands of weight, power and physical space. Increasing the number of beams per PAA results in higher cost and thermal load per unit. Therefore, there is a need to implement a PAA based efficient network design.

[0005] Synchronous links can be used to implement a backbone of a network. In such an approach, the nodes on the backbone communicate among themselves using long-duration synchronous links. In real world networks, e.g., military networks, a single network backbone is connected to a number of terminal nodes. The backbone to terminal links use short bursts of RF energy in a time division multiple access (TDMA) approach. This requires the backbone nodes to have two sets of PAAs, one set to cover 4π steradian sky for maintaining continuous links to other backbone nodes, and another set to cover 4π steradian sky for TDMA links connecting to the terminal nodes. Two full sets of antennas make additional demands in terms of space, power and weight to accommodate the extra set of antennas. For military aircraft, such demands are prohibitive. Further, the mounting of extra antennas requires additional aircraft fuselage penetrations. Thus, there is a need for a network with a single set of PAAs to communicate with both the backbone and terminal nodes in the network.

[0006] In a synchronous link based backbone network, a group of terminal nodes communicating with the TDMA links cannot function as an independent network, i.e., a peer-to-peer network, without connecting to some backbone node. In military environments, a backbone node may not be accessible in some critical circumstances, making the network unavailable. Hence, there is a need for a PAA based network which will allow formation of a network made exclusively of terminal nodes.

SUMMARY OF THE INVENTION

[0007] The present invention is directed to a TDMA network backbone. In a preferred form the network backbone is constructed from nodes having interconnecting links. The nodes have a set of receive antennas and a set of transmit antennas. The receive antennas are operated in accordance with receiving slots of a TDMA sequence and the transmit antennas are operated in accordance with transmitting slots of a TDMA sequence. A timing reference subsystem provides timing signals to the nodes. A node control is used to communicate information related to the nodes.

[0008] In one embodiment the transmit times are aligned to the boundaries of the transmit TDMA slots. In another embodiment the receive times are aligned to the boundaries of the TDMA receive slots. In one embodiment the receive and transmit antennas comprise phased array antennas.

[0009] In a preferred embodiment the links are formed by a given first node sending an order-wire message using the node control to a given second

node. The node control on the first node computes the time-of-flight (TOF) for transmission from the first node to the second node. The first node sends an order-wire message to the second node offering first transmit slots and first receive slots. The second node intersects the first transmit slots with second receive slots on the second node to determine an acceptable transmit list. In a similar manner, the second node determines an acceptable receive list. The second node sends an acceptance status message along with the acceptable transmit and receive lists if the lists are non-empty, otherwise the second node sends a rejection status message. The transmission link between the first and the second node is made operative by the first node once the acceptance status message is received.

[0010] In one embodiment the capacity of the transmission link is adjusted by either removing or adding the receive and transmit slots associated with the transmission link.

[0011] Further areas of applicability of the present invention will become apparent from the detailed description provided hereinafter. It should be understood that the detailed description and specific examples, while indicating the preferred embodiment of the invention, are intended for purposes of illustration only and are not intended to limited the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The present invention will become more fully understood from the detailed description and the accompanying drawings, wherein:

[0013] Figure 1 is a block diagram of a PAA based network;

[0014] Figure 2 is a shows a representation of a TDMA cycle;

[0015] Figure 3 illustrates a sending mechanism;

[0016] Figure 4 illustrates a receiving mechanism;

[0017] Figure 5 shows a spatial representation of a transmission;

[0018] Figure 6 illustrates a link formation process;

[0019] Figure 7 illustrates a receive block processing mechanism;

[0020] Figure 8 illustrates the calculating of send times;

[0021] Figure 9 illustrates a transmit block processing mechanism;

[0022] Figure 10 illustrates a user defined rule applied to physically acceptable lists;

[0023] Figure 11 illustrates non-overlapping bursts;

[0024] Figure 12 illustrates overlapping bursts;

[0025] Figures 13 and 14 show the plot of an average number of node hops per end-to-end route;

[0026] Figures 15 and 16 show an average link load;

[0027] Figures 17-19 show the fraction of individual overloaded links;

[0028] Figures 20 and 21 show an average processing load imposed on the nodes for a given amount of traffic; and

[0029] Figures 22 and 23 show an average end-to-end route latency.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0030] The following description of the preferred embodiment(s) is merely exemplary in nature and is in no way intended to limit the invention, its application, or uses.

[0031] Figure 1 shows a simplified block diagram of a network 10 in accordance with a preferred embodiment of the present invention. The network 10 consists of a plurality of nodes 12. The contents of one of the nodes 12a is also shown in Figure 1. All the nodes 12 have preferably similar configurations. Node 12a has at least a pair of PAAs 14. A subset of the total number of PAAs 14 on the node 12a functions as receiving PAAs, while the other subset of the total number of PAAs 14 functions as transmitting PAAs. For explanatory purposes, the receiving PAA will be denoted by reference numeral 14a and the transmitting PAA by reference numeral 14b. Also, while the following description may reference node 12a, it will be appreciated that each of the nodes 12 operate in identical fashion to node 12a.

[0032] Each of the nodes 12 includes a clock 16 functioning as a backup timing device to be used when a wide area timing reference 18 is unavailable. The wide area timing reference 18 provides timing signals for synchronizing inter-node communications. A variety of different types of wide

area timing references 18 can be used. For example, the network 10 preferably uses a global positioning system (GPS) as the wide area timing reference 18.

[0033] The backbone of the network 10 is a set of links among the nodes 12. For using such links any of the nodes 12 must find the other nodes 12 located within its line-of-sight and the transmit frequencies used by each of them. A node control 20, located on each of the nodes 12, sends information about the existence, location and transmit frequencies of the one node 12 to the other nodes 12. The node control 20 also receives similar information from the other nodes 12.

[0034] The node control 20 is used by the node 12 to send order-wire messages to the other nodes 12 requesting formation of a link. The node control 20 includes an algorithm to select one or more available frequency bands for transmission. For example, the node control 20 can use an algorithm that considers the bands used by the nearby nodes 12 and then selects the frequency band whose nearest user is farthest from the node 12 implementing the algorithm. Thereafter, the node control 20 uses a second algorithm to select one of the other nodes 12 in an attempt to form a link. For example, the second algorithm computes the range to each one of the other nodes 12 and selects the closest of those other nodes 12. Once the node control 20 selects a particular other node 12, the link formation process can be started. Those skilled in the art will appreciate that the invention is not limited by the type of the algorithm used for selecting the frequency bands for transmission.

[0035] Figure 2 shows a representation of a TDMA cycle 22. The TDMA cycle 22 consists of "S" number of time slots 24 numbered from 0 to S-1. The length of each slot 24 is set to ' Δt_{slot} ,' which is equal to the duration of an RF burst plus a small guard interval to accommodate timing discrepancies between the slots 24. The duration of the overall TDMA cycle 22 is given by the product $S \times \Delta t_{\text{slot}}$. Wide area timing reference 18 (as shown in Figure 1) is used to align the slots 24. For example, if the duration of Δt_{slot} is 100 microseconds and there are 100 slots 24, then the overall duration of the TDMA cycle 22 is 0.01 second.

[0036] The node 12a uses the node control 20 to determine its location. The location is determined within a fraction of the distance that an RF burst covers in duration length of the one slot 24. The given node 12a calculates the time-of-flight (TOF) to another node 12 after determining the location of the other node 12. As the locations for the nodes 12 are imprecisely known, each of the slots 24 includes a guard interval. The guard interval prevents successive RF bursts received by the adjoining slots 24 from overlapping each other. The length of the guard interval must be large if the locations of the nodes 12 are known less precisely. If the guard interval is made too long then the length of the slot 24 will also increase, resulting in a loss of efficiency.

[0037] Either the transmit times or the receive times are aligned on the boundaries of the slots 24, but not both. For example, if the transmit times are aligned on the boundaries of the slots 24, then every node 12 will transmit RF bursts at the same time, but the RF bursts would be received at different times depending on the time-of-flight (TOF). Similarly, the network 10 can be

configured so that the receive times are aligned on the boundaries of the slots 24. Those skilled in the art will readily appreciate that choosing either alignment is equally effective, and the invention is not limited in any manner by such a choice.

[0038] Figure 3 illustrates the sending protocol. The following discussion assumes that a choice has been made to align the receive times on the boundaries of the slots 24. A node "A" designated by reference numeral 26, transmits RF bursts to a node "B," designated by numeral 28, node C, designated by numeral 30, node D, designated by numeral 32 and node E, designated by numeral 34. Distances between the node A 26, which is transmitting, and the other nodes 12 varies depending on the TOF and is measured in the multiples of duration length of the slot 24. For example, node A 26 is at a distance of 10.25 slots from node B, indicating a TOF 10.25 times longer than the length of the single slot 24. Node A 26 starts transmitting to Node B 28 before a scheduled receive time for node B 28. The time by which transmission of node A 26 precedes the scheduled receive time for node B 28 is equal to the TOF between node A 26 and node B 28, which is 10.25 slots. Similarly, node A 26 will start transmitting to node C 30, node D 32 and node E 34 before the scheduled receive times for those slots. As the receive times are aligned to the slot boundaries, node B 28, node C 30, node D 32 and node E 34 receive RF bursts at the appropriate slot boundary depending upon their respective TOFs from node A 26.

[0039] Figure 4 illustrates the receiving protocol. Here, node B 28, node C 30, node D 32 and node E 34 transmit towards node A 26. Node A 26 expects to receive RF bursts at well aligned receive times. Node B 28, node C 30, node D 32 and node E 34 transmit towards node A 26 typically at non-aligned times; these nodes start transmitting at times equal to the TOF before the scheduled receive slot for node A 26.

[0040] Figure 5 shows a spatial representation of a burst received by node A 26. Node B 28, node C 30, node D 32 and node E 34 transmit towards node A 26 at varied times depending upon the TOF between each particular one of these nodes and the receiving node A 26. The bursts are spaced to arrive at successive and non-overlapping times which are aligned to the specific time slots 24.

[0041] Figure 6 illustrates an offered transmit slots block 36 and an offered receive slots block 38.

[0042] Referring to Figures 7-9, the operation of the receive block processing protocol, together with the calculation of send times and the operation of the transmit block processing protocol will be described. For illustrating the link formation process, node A 26 and node B 28 (refer to figure 5) are considered as examples described below. Node A 26 identifies the transmitting PAA 14b that points closest towards node B 28 using the locations of node A 26 and node B 28 along with the orientation of node A 26. Node A 26 calculates the TOF using the locations of nodes A 26 and B 28.

[0043] Node A 26 calculates the start times for sending bursts from node A 26 in order to transmit to node B 28 by subtracting the TOF from the starting time of the slot 24. Node A 26 thus builds a list of start times for sending bursts from node A 26 to node B 28. Node A 26 determines which of the available intervals on the array of transmit PAA 14b are previously unassigned and hence available for transmitting bursts to node B 28. Node A 26 puts the available slot 24 indicator on the offered transmit slots block 36 (Figure 7) to node B 28 if for each interval at least one of the beams of transmit PAA 14b is not already assigned to any of the links with the other nodes 12. The available slot 24 indicator corresponds to the receive slot 24 during which node B 28 would receive the possible burst.

[0044] Node A 26 determines the slots 24 on the receive array of its PAAs which are available for receiving RF bursts from node B 28. Node A 26 determines, for each of the 'S' number of the slots 24, whether at least one of the beams of the receiving PAA 14a is not previously assigned to links with the other nodes 12, then the node A 26 puts the identifier for that slot 24 on the offered receive slots block 38 (Figure 8) to be offered to node B 28.

[0045] Node A 26 sends an orderwire message to node B 28 tendering offered transmit slots block 36 and the offered receive slots block 38. The offered transmit slots block 36 indicates times at which node A 26 could transmit bursts from node A 26. The offered receive slots block 38 indicates the times at which node A 26 has the capacity to receive bursts from the node B 28.

[0046] Node B 28 then analyzes the orderwire request sent by node A 26. Node B 28 determines which of its receiving PAA 14a is pointing closest to node A 26, and which of the slots 24 are available for its receiving PAA 14a to receive bursts from node A 26.

[0047] With further reference to Figure 7, node B 28 intersects the available receive slots 24 and the available transmit slots 24 from node A 26. First, the intersection process for receiving slots 24 will be discussed. For illustration purposes, the receiving PAAs 14a located on node B 28 have two beams, i.e., a first receive beam 40 and a second receive beam 42. Node B 28 intersects the offered transmit slots block 36 received from node A 26 to the slots 24 for the first receive beam 40 and a second receive beam 42. Node B 28 accepts those transmit slots from offered transmit slots block 36 for which the receive slot 24 is available on either the first receive beam 40 or a second receive beam 42 or both of them, and such acceptable slots 24 are put on a physically acceptable receive list 46. Those slots on the offered transmit slots block 36 for which both the first receive beam 40 and the second receive beam 42 are unavailable are ignored as unavailable receive slots 44 and hence are not reflected on the physically acceptable receive list 46.

[0048] Secondly, node B 28 carries out the intersection process for the transmitting slots 24, node B 28 determines which of its transmitting PAAs 14b is closest to node A 26 and then computes the TOF from itself to node A 26. Referring to Figure 8, node B 28 computes start times corresponding to each of the slots 24 on the offered receive slots block 38 to create a send times block 48.

Node B intersects the send times block 48 with a first transmit beam 50 and second transmit beam 52, as indicated in Figure 9. Node B 28 accepts those slots 24 from send times block 48 for which the transmit slot 24 is available on either the first transmit beam 50 or second transmit beam 52, or on both of them, and such acceptable receive slots 24 are put on a physically acceptable transmit list 54. For those slots on the send times block 48 for which both the first transmit beam 50 and second transmit beam 52 are unavailable, they are ignored as unavailable transmit slots 56 and not reflected on the physically acceptable transmit list 54.

[0049] Figure 10 illustrates a user defined rule applied to physically acceptable lists. Node B applies any user defined rules to physically acceptable receive list 46 and the physically acceptable transmit list 54. The user defined rule illustrated in figure 10 is termed as "avoid consecutive bursts" , which is used to edit the physically acceptable transmit list 54. As shown in Figure 10, the physically acceptable transmit list 54 is edited to remove consecutive bursts as per the user defined rule. Similar rules can be used for the physically acceptable receive list 46.

[0050] Node B 28 sends an acceptance message to node A 26 with the selected receive and transmit slots 24. Node B 28 sends the acceptance message only if at least one slot 24 remains in each of the receive and transmit lists and the link meets any user specified criteria, otherwise it sends a rejection message to node A 26. If the proposed link was accepted by node B 28, then node A 26 and node B 28 point their respective receiving PAAs 14a and

transmitting PAAs 14b at each other at agreed upon times and also begin transmitting to each other at agreed on upon times. Thus, a link is formed between node A 26 and node B 28. In order to maintain the link, node A 26 and node B 28 intermittently receive updated location information from each other, based on which node A 26 and node B 28 adjust their beam pointing and burst timing.

[0051] In case node A 26 does not receive a response from node B 28, then node A 26 repeats the process a given number of times before abandoning the attempt to form a link with node B 28. After abandoning attempts to form a link with node B 28, node A 26 attempts to form links with the other nodes 12. If node B 28 does not receive any scheduled bursts within a user-defined interval after sending an acceptance message to node A 26, then node B 28 retransmits the acceptance message a user-defined number of times. Node B 28 abandons the link if node A 26 fails to transmit even after node B 28 retransmits the acceptance message a user-defined number of times.

[0052] The TDMA backbone approach allows the nodes 12 to adjust capacity of each link to match the demand for data traffic. The link capacity is effectively boosted to carry more data over an existing link. For illustrating the process of increasing the capacity of a link, the link between node A 26 and node B 28 is considered. Node A 26 estimates the number of additional slots 24 which need to be added to the TDMA cycle 22 for additional capacity. For example, if node A 26 needs to transmit an additional 1 million bits per second and each burst carries 10,000 bits, then node A 26 would need to add 100 bursts

per second to the link. If the duration of the TDMA cycle 22 is 100 milliseconds, then there are ten TDMA cycles 22 per second. Hence, to add 100 bursts per second would mean adding 10 bursts per TDMA cycle 22. Thus, the number of slots 24 that need to be added for additional capacity is 10 more slots 24 per TDMA cycle 22.

[0053] Node A 26 subtracts the TOF from the starting time for each slot 24 in the TDMA cycle 22 using the already known TOF for transmissions to node B 28. The result is a send times block 48. Using the already known physically acceptable transmit list 54, node A 26 determines which transmit intervals are available and accordingly updates the physically acceptable transmit list 54. Node A 26 sends a request for additional transmit capacity to node B 28, along with the number of additional slots that node A 26 would like to add and the list of physically acceptable transmit intervals. On receiving the message for additional capacity from node A 26, node B 28 determines which of the slots 24 are available on its receiving PAAs 14a to receive bursts from node A 26. Node B 28 computes the intersection of the available slots 24 on its receiving PAAs 14a and the slots 24 offered by node A 26. The result of such intersection is put on the physically acceptable receive list 46 for node B32. Node B 28 removes slots from the physically acceptable receive list 46 until the number of the slots 24 is not greater than the number of additional slots 24 required by node A 26.

[0054] Node B 28 sends a response to node A 26 with the selected additional slots 24 (possibly none). If node B 28 has responded with at least one additional slot 24, then node B 28 begins pointing its receiving PAAs 14a during

the additional slot 24, and node A 26 begins transmitting to node B 28 at the agreed upon times. Thus, the additional capacity becomes operational. If node B 28 accepts a lesser number of slots 24 than those requested by node A 26, then node A 26 routes excess traffic via other links.

[0055] At certain times link capacity will be adjusted by removing excess capacity from the link. For illustrating the excess capacity reduction, the link between node A 26 and node B 28 is considered. Node A 26 estimates the number of receive slots 24 and the number of transmit slots 24 that can be removed from each of the TDMA cycles 22. For example, in one situation each burst carries 10,000 bits, and node A 26 estimates that the capacity of the link exceeds the demand by one million bits per second for the data transmitted to node B 28 and two million bits per second for the data received from node B 28. Node A 26 would then remove 100 bursts per second from node A 26 to node B 28 part of the link, and 200 bursts per second from the node B 28 to the node A 26 part of the link. If the duration of the TDMA cycle 22 is 100 milliseconds, then there will be 10 TDMA cycles 22 per second, which means removing 100 bursts per second equates to removing 10 bursts per TDMA cycle 22.

[0056] Node A 26 applies user defined criteria to select the slots 24 to be removed from the physically acceptable receive list 46 and the physically acceptable transmit list 54. For example, the user defined criteria could be to remove the slots 24 that are needed to meet demand on another link on the same PAA. Alternatively, the user defined criteria could be to specify a minimum number of the slots 24 per link, even when the link is lightly loaded, to maintain

low latency. Node A 26 sends a message to node B 28 specifying which of the transmit and receive slots 24 are to be removed from the link. Node B 28 removes the specified slots 24 from the physically acceptable receive list 46 and the physically acceptable transmit list 54. Node B 28 then sends an acknowledgement to node A 26. Both node A 26 and node B 28 cease using the specified transmit slots 24, as well as the specified receive slots 24. Thus, node A 26 and node B 28 both are free to assign the removed transmit slots 24 and receive slots 24 to other links.

[0057] Network 10 typically operates in a mobile environment. Some of the nodes 12 will be in motion relative to other nodes 12. For example, as node A 26 and node B 28 move through the atmosphere, it is possible that node B 28 will move out from the field-of-regard of a first PAA 14 on node A 26 and into the field-of-regard of the second PAA 14 also located on node A 26. A conflict could possibly occur because the slots currently used for the link on first PAA 14 could already be assigned to other links on the second PAA 14.

[0058] Node A 26 and node B 28 take steps to avoid conflicts when a link transfers from one PAA 14 to another PAA 14. Node A 26 continuously tracks the angular position and angular rate of each link relative to the field-of-regard of a given PAA 14. Node A 26 continuously predicts when each link is likely to cross the edge of the field-of-regard for a given PAA 14. Node A 26 considers its own rotation and/or the position of node B 28, and estimates that typically within a few TDMA cycles 22 that node B 28 will be in the field-of-regard of the second PAA 14. Node A 26 identifies the available slots on the second

PAA 14 and sends a list of the available slots 24 to node B 28 along with the estimated time at which node B 28 will be in the field-of-regard of the second PAA 14. Node B 28 identifies which of the offered slots 24 it can accept, using the intersection process described above. Node B 28 sends the list of accepted slots 24 to the node A 26. Node A 26 and node B 28 then start transmitting at agreed upon times using the newly agreed upon slots 24.

[0059] Figure 11 illustrates non-overlapping bursts from node A 26 to node B 28, node C 30, node D 32 and node E 34. As the nodes move, the times of flight from node A 26 to the other nodes may vary at differing rates leading to overlap in the transmit intervals on different links. Figure 12 illustrates overlapping bursts. Node A 26 must eliminate any overlapping burst, but it can maintain the link to node C 30 by adding a burst at any of the slots 24 filled with cross-hatchings in Figure 12.

[0060] A special situation results when previously overlapping blocks become available "just-in-time." For example, referring to Figure 11, the fourth block from the left in the link to node C 30 overlaps a block in transmission from node A 26 to node B 28. The fourth block from the left is therefore not available for transmissions from node A 26 to node C 30. But just when the leftmost block on the link to node C 30 (indicated by an "X") begins to overlap another transmission from node A 26 to node B 28, the fourth block from the left on the link to node C 30 which was previously overlapping ceases to overlap the transmission from the node A 26 to the node B 28. Thus, certain slots 24 which

were previously locked in an overlap may become available “just-in-time” as the overlap pattern changes.

[0061] To overcome special situations resulting from changes in distances between nodes, the transmit intervals need to be reassigned. Node A 26 tracks the timing of transmit bursts in each link relative to bursts in other links on the same transmitting PAA 14b, while continuously predicting the likelihood of the bursts overlapping each other. Node A 26 estimates that the relative motion of node B 28 and node C 30 will put one or more bursts from node A 26 to node B 28 into conflict with the bursts from node A 26 to node C 30 within a span of a few TDMA cycles 22. Node A 26 identifies available slots on the array of the transmitting PAAs 14b which will become available when the conflict starts. Node A 26 sends the list of slots 24 that must be deleted to avoid overlaps and a list of available slots including the “just-in-time” available slots 24 to node C 30, along with the estimated time at which the changes must be made. Node C 30 identifies the slots 24 which it can accept using the intersection methods described above and sends the list of acceptable slots 24, if any, to node A 26. At agreed upon times, node A 26 and node C 30 cease to use the conflict causing slots 24 and start using the newly agreed upon slots 24.

[0062] Networks having a TDMA backbone generally perform better than conventional networks based on the synchronous approach. Figures 13-23 show graphs of simulations comparing the invention with the conventional synchronous approach. Nodes 12 in the network analyzed in the Figures 13-23 were distributed throughout a 2000 km x 2000 km theatre of war, and the

altitudes varied from sea level up to 10 km above sea level. All traffic was routed along the shortest path with a minimum number of hops. The 'TAP' acronym in the legend of these figures refers to 'TMIG Access Point' (TMIG stands for Theater Mobile Information Grid), which is equivalent to node 12, and further discussion will assume that they are similar.

[0063] Figures 13 and 14 show the plot of the average number of transmission hops per end-to-end route for 50 and 125 nodes, respectively. A network with fewer hops tends to be more efficient. As illustrated, the TDMA backbone of the invention yields fewer hops regardless of the number of nodes in the network. The effectiveness of the invention increases as the number of nodes increases.

[0064] Figures 15 and 16 show the average link load, which is the amount of traffic carried over each link divided by the capacity of that link. The synchronous approach provides lower average loads for networks with few nodes. The invention and the synchronous approach provide comparable average loads at 75 to 100 nodes. The invention, at 125 or more nodes, imposes only about 85 percent of the load of a synchronous backbone.

[0065] Figures 17-19 show the fraction of individual overloaded links, i.e., the links for which the amount of traffic over the link exceeds the capacity of the link. The invention and the synchronous approach provide similar results for networks with 50 nodes. For 75 or more nodes the invention is superior in almost all cases over the synchronous approach. With 125 nodes, the number of overloaded links for the invention is only 29% as great as for the synchronous

backbone. In real networks, as opposed to present simulations with least hops, routers tend to direct traffic around overloaded links. Synchronous links in real networks would have lesser overloaded links, but will have more hops. Thus, synchronous approach in real networks would suffer from increases in link loading, processing load and latency.

[0066] Figures 20 and 21 show the average processing load imposed on the nodes for a given amount of traffic. The invention provides a lower processing load in all cases and at all network sizes. As the number of nodes increases the invention provides even better processing load. For example, at 125 nodes the invention imposes only 53% of the processing load imposed by a synchronous backbone.

[0067] Figures 22 and 23 show the average end-to-end route latency. The invention provides lower latency in all cases and all network sizes. The invention provides even better performance as the number of nodes increases. For example, at 125 nodes, the latency with the invention is only 76% of the latency with a synchronous backbone.

[0068] As compared to conventional approaches, the present invention provides a lower cost mobile network by permitting a lower data rate per beam (i.e. lower modem speeds) and lower processing speed in each node (slower decryption, error correcting, and routing speed) to achieve the same end-to-end data rate. The invention also provides lower end-to-end latency than conventional approaches. The invention furthermore permits multiple nodes 12 to transmit simultaneously and does not require burst durations longer than the

TOF. In addition, the invention permits a single set of PAAs covering 4π steradian sky on a network backbone node to provide network backbone communications together with TDMA links to terminal nodes.

[0069] The description of the invention is merely exemplary in nature and, thus, variations that do not depart from the gist of the invention are intended to be within the scope of the invention. Such variations are not to be regarded as a departure from the spirit and scope of the invention.